

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

- 1 1. (Currently Amended) A method of processing traffic received from
2 an InfiniBand node via a first queue pair, comprising:
3 selecting a traffic entry in an InfiniBand receive queue, wherein said traffic
4 entry comprises one of:
5 a Send command comprising an encapsulated communication;
6 a Send command comprising an RDMA Read descriptor; and
7 an RDMA Read response comprising a response to an RDMA
8 Read request;
9 if said selected traffic entry comprises a Send command comprising an
10 RDMA Read descriptor:
11 issuing a first RDMA Read request to retrieve one or more portions
12 of a communication described by said RDMA Read descriptor;
13 in a linked list corresponding to the first queue pair, adding an
14 entry corresponding to said first RDMA Read request, said entry
15 comprising a first sequence number and a last sequence number that
16 identify a range of sequence numbers associated with expected responses
17 to said first RDMA Read request and optionally comprising a sequence
18 number of the most recently received response or a link to the next entry of
19 in the linked list~~identifying a range of sequence numbers associated with~~
20 ~~expected responses to said first RDMA Read request;~~ and
21 in a retry queue, adding an entry corresponding to said first RDMA

22 Read request; and
23 if said selected traffic entry comprises an RDMA Read response to said
24 first RDMA Read request:
25 identifying a sequence number associated with said RDMA Read
26 response;
27 comparing said sequence number to said range of sequence
28 numbers;
29 storing said one or more portions of said described communication
30 to facilitate assembly of said described communication in said queue; and
31 if said sequence number matches a final sequence number in said
32 range, retiring in said retry queue said entry corresponding to said first
33 RDMA Read request.

1 2. (Original) The method of claim 1, further comprising:
2 forwarding a communication associated with said selected traffic entry, for
3 transmission on an external communication link, wherein said communication is
4 one of:
5 said encapsulated communication; and
6 said described communication, after said described communication
7 is assembled.

1 3. (Original) The method of claim 1, further comprising, if said
2 selected traffic entry comprises an RDMA Read response to said first RDMA
3 Read request:
4 if said sequence number does not match said final sequence number,
5 updating said entry in said linked list to include said sequence number.

1 4. (Original) The method of claim 1, further comprising:

2 maintaining a single memory structure comprising multiple linked list,
3 including said linked list;
4 wherein each linked list stores entries associated with RDMA Read
5 requests for a different InfiniBand queue pair.

1 5. (Original) The method of claim 1, further comprising:
2 maintaining a single memory structure for queuing InfiniBand traffic
3 received via multiple virtual lanes and multiple queue pairs, said single memory
4 structure comprising said queue.

1 6. (Original) The method of claim 5, wherein said queue comprises a
2 linked list of memory buffers within said single memory structure.

1 7. (Original) The method of claim 1, further comprising:
2 maintaining a head pointer configured to identify a head of said linked list;
3 and
4 maintaining a tail pointer configured to identify a tail of said linked list.

1 8. (Currently Amended) The method of claim 1, further comprising:
2 maintaining a head pointer configured to identify a head of said queue;
3 maintaining a tail pointer configured to identify a tail of said queue; and
4 maintaining a next traffic entry pointer configured to identify a next entry
5 in said queue to be processed ~~after said forwarding~~.

1 9. (Original) The method of claim 8, wherein said tail pointer is
2 configured to identify where in said queue a next traffic entry is to be queued.

1 10. (Original) The method of claim 1, further comprising, if said

2 selected traffic entry comprises an RDMA Read descriptor:
3 appending space to a head of said queue;
4 wherein said described communication is assembled in said appended
5 space.

1 11. (Original) The method of claim 1, further comprising, if said
2 selected traffic entry comprises an RDMA Read response to said first RDMA
3 Read request:
4 dropping an RDMA Read response received out of order; and
5 requesting a retry of said first RDMA Read request.

1 12. (Currently Amended) A computer readable medium storing
2 instructions that, when executed by a computer, cause the computer to perform a
3 method of processing traffic received from an InfiniBand node via a first queue
4 pair, the method comprising:
5 selecting a traffic entry in an InfiniBand receive queue, wherein said traffic
6 entry comprises one of:
7 a Send command comprising an encapsulated communication;
8 a Send command comprising an RDMA Read descriptor; and
9 an RDMA Read response comprising a response to an RDMA
10 Read request;
11 if said selected traffic entry comprises a Send command comprising an
12 RDMA Read descriptor:
13 issuing a first RDMA Read request to retrieve one or more portions
14 of a communication described by said RDMA Read descriptor;
15 in a linked list corresponding to the first queue pair, adding an
16 entry corresponding to said first RDMA Read request, said entry
17 comprising a first and a last sequence numbers that identify a range of

18 sequence numbers associated with expected responses to said first RDMA
19 Read request and optionally comprising a sequence number of the most
20 recently received response or a link to the next entry of in the linked list
21 ~~identifying a range of sequence numbers associated with expected~~
22 ~~responses to said first RDMA Read request; and~~
23 in a retry queue, adding an entry corresponding to said first RDMA
24 Read request; and
25 if said selected traffic entry comprises an RDMA Read response to said
26 first RDMA Read request:
27 identifying a sequence number associated with said RDMA Read
28 response;
29 comparing said sequence number to said range of sequence
30 numbers;
31 storing said one or more portions of said described communication
32 to facilitate assembly of said described communication in said queue; and
33 if said sequence number matches a final sequence number in said
34 range, retiring in said retry queue said entry corresponding to said first
35 RDMA Read request.

1 13. (Currently Amended) A method of tracking responses to an RDMA
2 Read operation, the method comprising:
3 issuing an RDMA Read on a first communication connection;
4 identifying a range of sequence numbers to be associated with responses to
5 the RDMA Read;
6 adding an entry to a first linked list corresponding to the first
7 communication connection, said entry comprising ~~configured to identify:~~
8 a first and a last sequence numbers that identify said range of
9 sequence numbers; ~~and~~

10 a latest sequence number received in said range of sequence
11 numbers;
12 optionally a link to the next entry of in the linked list;
13 receiving a first RDMA Read response;
14 determining whether a first sequence number associated with the first
15 RDMA Read response matches a last sequence number in said range of sequence
16 numbers; and
17 if said first sequence number does not match said last sequence number,
18 updating said latest sequence number to match said first sequence number.

1 14. (Original) The method of claim 13, further comprising:
2 if said first sequence number matches said last sequence number, retiring
3 an entry in a retry queue corresponding to the RDMA Read.

1 15. (Original) The method of claim 13, wherein the first
2 communication connection is an InfiniBand queue pair.

1 16. (Original) The method of claim 15, wherein:
2 said issuing is performed by an InfiniBand transmit module; and
3 said adding, said determining and said updating are performed by an
4 InfiniBand receive module;
5 the method further comprising:
6 at the InfiniBand transmit module, retrying the RDMA Read if an RDMA
7 Read response associated with said range of sequence numbers is received out of
8 order.

1 17. (Original) The method of claim 15, wherein:
2 said issuing is performed by an InfiniBand transmit module; and

3 said adding, said determining and said updating are performed by an
4 InfiniBand receive module;
5 the method further comprising, at the InfiniBand transmit module:
6 maintaining a retry queue for tracking RDMA Reads that have not
7 yet completed; and
8 retiring an entry in said retry queue corresponding to the RDMA
9 Read if RDMA Read responses corresponding to said range of sequence
10 numbers are received in order.

1 18. (Original) The method of claim 17, further comprising, at the
2 InfiniBand transmit module:
3 retrying the RDMA Read if RDMA Read responses corresponding to one
4 or more of said range of sequence numbers are received out of order.

1 19. (Original) The method of claim 13, wherein said identifying
2 comprises:
3 dividing an amount of data to be received in response to the RDMA Read
4 by a maximum transfer unit in effect for the first communication connection.

1 20. (Original) The method of claim 13, further comprising:
2 maintaining a single memory structure comprising multiple linked lists
3 corresponding to multiple communication connections, including said first linked
4 list corresponding to the first communication connection.

1 21. (Original) The method of claim 20, further comprising:
2 for each of the multiple communication connections, including the first
3 communication connection, maintaining pointers to the first entry and the last
4 entry in the corresponding linked list.

1 22. (Original) An apparatus for queuing multiple types of receive
2 traffic in a communication interface, comprising:
3 a queue for queuing multiple types of receive traffic associated with
4 communications to be transmitted from the communication interface;
5 a head pointer configured to identify a head of said queue;
6 a tail pointer configured to identify a tail of said queue, wherein said
7 traffic commands are enqueued at said tail;
8 a next entry pointer configured to identify a next entry in said queue to be
9 processed; and
10 a linked list, wherein each entry in said linked list corresponds to an
11 RDMA Read request issued by the communication interface, and is configured to
12 identify a range of sequence numbers associated with expected responses to the
13 RDMA Read request.

1 23. (Original) The apparatus of claim 22, wherein each entry said
2 linked list is further configured to identify a sequence number of a most recently
3 received response to the RDMA Read request.

1 24. (Original) The apparatus of claim 22, wherein the linked list is one
2 of multiple linked lists, each said linked list corresponding to a separate
3 InfiniBand queue pair.

1 25. (Original) The apparatus of claim 22, further comprising:
2 a retry queue, wherein a retry entry is added to said retry queue for each
3 RDMA Read request issued by the communication interface;
4 wherein a first retry entry in said retry queue corresponding to a first
5 RDMA Read request is retired when said expected responses to the first RDMA

6 Read request are received.

1 26. (Original) The apparatus of claim 22, further comprising:
2 a memory configured to store pointers to a first entry and a last entry in
3 said linked list.

1 27. (Original) The apparatus of claim 22, wherein said queue
2 comprises an assembly area for assembling a communication associated with a
3 first type of receive traffic.

1 28. (Original) The apparatus of claim 27, wherein said assembly area
2 comprises a portion of said queue delimited by said head pointer and said next
3 entry pointer.

1 29. (Original) The apparatus of claim 27, wherein said first type of
2 receive traffic is an InfiniBand RDMA Read command comprising a set of
3 RDMA read descriptors configured to identify the communication associated with
4 said first type of receive traffic.

1 30. (Original) The apparatus of claim 29, wherein a second type of
2 receive traffic is an InfiniBand Send command configured to encapsulate the
3 communication associated with said second type of receive traffic command.

1 31. (Original) The apparatus of claim 27, wherein:
2 said first type of receive traffic comprises a set of descriptors, wherein
3 each said descriptor is configured to describe a portion of the communication
4 associated with said first type of receive traffic; and
5 the apparatus is configured to issue read requests to retrieve the portions of

6 the communication described by the set of descriptors and assemble said portions
7 in said assembly area.

1 32. (Original) The apparatus of claim 22, further comprising:
2 a transmit module configured to transmit the communications associated
3 with said receive traffic;
4 wherein each communication associated with receive traffic is forwarded
5 from said queue to said transmit module after the communication is determined to
6 be complete.

1 33. (Original) The apparatus of claim 32, wherein a communication is
2 forwarded from said queue to said transmit module by passing to the transmit
3 module a set of pointers delimiting the communication within said queue.

1 34. (Original) The apparatus of claim 22, wherein said queue
2 comprises a linked list of buffers within a memory structure configured to queue
3 receive traffic for multiple communication connections.

1 35. (Currently Amended) A communication interface for tracking
2 responses to an InfiniBand RDMA Read request, comprising:
3 for each of one or more active InfiniBand queue pairs, a corresponding
4 linked list, wherein each entry in said linked list is configured to include:
5 a first and a last sequence numbers that identify a range of
6 sequence numbers associated with expected responses to an RDMA Read
7 request issued on the corresponding queue pair by the communication
8 interface; ~~and~~
9 a previous sequence number, wherein said previous sequence
10 number is a sequence number associated with a most recently received

11 response to the RDMA Read request; and
12 optionally a link to the next entry of in said linked list; and
13 for each of the linked lists, pointers to a first entry and a last entry in said
14 linked list.

1 36. (Original) The communication interface of claim 35, further
2 comprising:
3 a retry queue configured to queue retry entries corresponding to RDMA
4 Read requests issued by the communication interface;
5 wherein an retry entry in said retry queue is retired when a final response
6 to a corresponding RDMA Read request is received, said final response being
7 identified by a final sequence number in said range of sequence numbers.

1 37. (Original) The communication interface of claim 35, further
2 comprising:
3 a transmit module configured to:
4 issue a first RDMA Read request on a first queue pair; and
5 calculate said range of sequence numbers associated with said
6 expected responses to the first RDMA Read request; and
7 a receive module configured to add an entry, corresponding to the first
8 RDMA Read request, to said corresponding first linked list.

1 38. (Original) The communication interface of claim 37, wherein said
2 receive module is further configured to:
3 determine a sequence number of a response to the first RDMA Read
4 request; and
5 determine if said sequence number matches a final sequence number in
6 said range of sequence numbers associated with expected responses to the first

7 RDMA Read request.

1 39. (Original) The communication interface of claim 38, wherein said
2 receive module is further configured to:
3 determine if said sequence number is out of order.